CHAPTER

4

SERVER POWER AND RACK INSTALLATION

After reading this chapter and completing the exercises, you will be able to:

- Identify features of server power supply
- Correctly implement an uninterruptible power supply (UPS)
- Plan optimum placement of equipment in a server rack
- Configure a keyboard, video, mouse (KVM) console
- List tips for installing equipment in racks

Power is an obvious requirement for a server, and this chapter shows you how server-level power supplies differ from standard desktop workstation power supplies. While a server power supply typically has a respectable mean time between failure (MTBF) of around two million hours, power to the office building does not. Many factors can cause power to fail to the server room, and you should prepare server room equipment for such an event with one or more devices that can temporarily provide power.

The byproduct of powering the server is the generation of heat. Fans are the primary method of protecting server components from overheating. It's almost impossible to install too many fans in the server, rack, or cabinet, and this chapter shows you various types of fans and how to strategically place them. This is especially important in the rack, where multiple server devices compound heating issues. Installing server equipment in the rack is a strategic process that requires planning in terms of power, heat, and weight distribution. In fact, the server rack often has so many installed servers that it is impractical to attach a separate keyboard, monitor, and mouse to each server. In that case, you will need to use a single keyboard, monitor, and mouse to switch between numerous servers.

Supplying Power to the Server

Obviously, without power you don't have a working server. Selecting a server with the type of power supply you need and ensuring its continued operation through the use of an uninterruptible power supply (UPS) is the first step in ensuring solid uptime percentages.

Power Supply

On the most basic level, an entry-level server probably has one power supply of at least 330 watts (W), whereas most workstations are probably as low as 145 W. A high-end, 8-way server probably provides between two and four power supplies rated between 375 W and 750 W each. The power supply, also known as the **power supply unit (PSU)**, attaches to the server chassis. Low-end servers and high-end workstations might not have hot-swappable and/or redundant power supplies. This means that in order to replace the power supply, you must remove the server case cover and disconnect power supply connectors to the motherboard, hard disks, floppy disks, and CD-ROMs before removing the screws attaching the PSU to the chassis and lifting out the power supply. This type of power supply installation is no different from the procedure on a typical desktop workstation because the chassis includes space for only a single, non-hot-swappable PSU. Well-configured servers, on the other hand, offer at least two hot-swappable power supplies. You can replace one of the power supplies without turning off the server (see Figure 4-1).

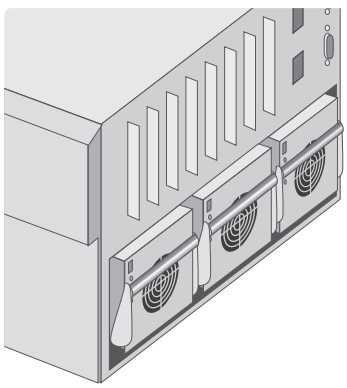

Figure 4-1 Redundant power supply units

Typically the server can operate acceptably with only a single PSU. The second power supply provides load balancing to reduce demands placed on a single PSU, as well as failover in case one of the power supplies fails. Often, a server includes three power supplies—two to provide continuous power and a third on standby in case one of the first two fails. Server management software and warning lights on the server should alert you to failed

or unstable power supplies, usually indicated by a fan that does not spin at appropriate levels (see Figure 4-2).

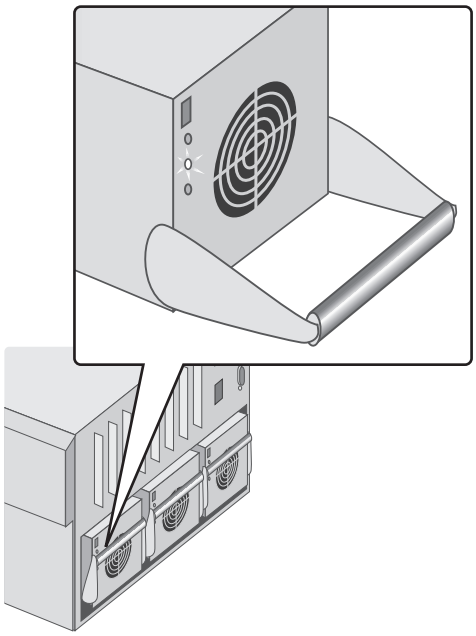

Figure 4-2 Most PSUs include a warning light that notifies you of problems

If one PSU fails, a standby comes online automatically and without interruption in service. (Some server configurations such as the Intel AD450NX server platform include support for an optional fourth PSU.) Because most networks require server availability 24/7, PSU failover is critical to a true server-level configuration. Also, the hot-swappable power supply does not require separate connectors for hard disks, CD-ROMs, the motherboard, and so forth. A hot-swappable PSU has sockets that plug directly into the power system of the server, which supplies power cables to server components. The power cord does not attach directly to the PSU (which would require a separate power cord for each PSU). Instead, a single power inlet serves all power supplies. Specific steps to replace a power supply are given in Chapter 6.

Many servers offer **N+1** expandability for critical components, particularly the processor and the PSU. "N" is a variable that refers to the quantity of a given component installed in a system, such as two power supplies. The "+1" refers to a spare component. For example, a server with three power supplies might be referred to as 2+1, in which two power supplies provide ongoing power while an additional power supply provides redundancy. N+1 can also refer to a chassis designed with space to accommodate additional components.

Each PSU requires a large fan to dissipate the heat generated by the power supply. While desktop workstation power supplies also have a fan, the server PSU often uses variable fan speed, which increases or decreases based upon the amount of heat the PSU thermistor detects. A **thermistor** is an internal thermostat that increases or decreases fan speed based on heat levels.

Calculating Server Power Requirements

Before calculating power requirements, make sure that the power from the building to the server room is sufficient to service your equipment. Although most building outlets are 110 volts (V), like household power outlets, server rooms usually also have 208/220 V outlets to accommodate the high power demands of larger servers with multiple power supplies and racks full of equipment. As part of your server room design, make sure the electrical engineers provide plenty of 208 V outlets for present and future needs. In a rack configuration, a **power distribution unit** (**PDU**), similar in function to a household "power strip" but with much higher capacity, often plugs into the 208 V outlet and supplies power to internal rack components (see more about racks later in this chapter). Server power supplies automatically detect the voltage of the power source and adjust as needed.

When a high-powered server or rack power is turned on, a sudden, temporary surge of power to the system takes place (known as **inrush power**). To account for this, be sure that the electrical engineer not only knows the amperage requirements for equipment that is up and running, but also accounts for inrush power, usually at least 20 amps (A) per PDU. Otherwise, you won't be able to power up a rack of equipment all at once without tripping the breaker.

Calculating the power supply needs for a server requires you to know how much power the motherboard, processor, internal adapters, and peripherals require. The power supply in the server can probably handle additional components without any problems; however, if you fill all expansion slots and drive bays, you might exceed power supply ratings. Though some components might list power requirements on the device or with its documentation, power requirements for other components might be difficult to locate. In that case, I recommend that you err on the safe side and calculate based on the maximum wattage allowed for a given type of device. Table 4-1 is a starting point.

Some devices might list volts, watts, amps, or combinations thereof. Use the information you gather to calculate power according to the following formula:

$$\text{watts} = \text{volts} \times \text{amps}$$

(This formula involves slight rounding, but should suffice for calculating general power requirements.) For example, a PCI slot requiring 5 V of 5 A current would require 25 W.

Computer components require positive power voltages in +3.3, +5, and +12 V. For example, a hard disk requires +12 V, and the processor usually requires +3.3 V. You might also find a negative power voltage of –5 V for backward compatibility with the ISA bus. The motherboard can be designed to supply negative voltages if the power supply does not. Assume all voltages in this book to be positive unless stated otherwise.

Table 4-1 Approximate Wattage Requirements

Component	Wattage Requirement
ATX motherboard (without CPU or RAM)*	30 W
RAM (approx. 10 W per 128 MB)	40 W for 512 MB
Pentium III 750 MHz	25 W
Floppy drive	5 W
IDE 50X CD or 10X DVD	25 W
4X AGP	30 W
PCI Card (5 W each)	30 W for six cards
IDE 5400 RPM drive**	10 W
IDE 7200 RPM drive	15 W
SCSI 7200 RPM drive	25 W
SCSI 10000 RPM drive	40 W

* 30 W represents a single, basic motherboard. Motherboard power requirements vary greatly depending on whether the server uses riser or mezzanine boards, which in turn require additional power.

** Hard drives require much more power during the spinup phase: 7200 RPM IDE drives require up to 30 W, and 10,000 RPM drives up to 40 W. Power requirements listed include **drive logic**, which is the circuitry included in the floppy or hard drive that interfaces with the disk controller. Many current hard drives are much more efficient in power requirements than the drives listed here.

Add the total power requirements for all server components, and subtract the total number from the power rating for the power supply. You should have plenty of power to spare, preferably about 6%, if you want the system to be as reliable as possible. If available power is marginal, you might consider moving certain components to other servers if possible, or upgrading the power supply.

Older computers use a paddle switch located on the power supply to turn on the computer. More recent servers and workstations use a remote power switch, which runs cables from the power source connecting leads to connectors on the switch. Be sure that if you are working inside the case with these wires, you disconnect the power cable first because the wires carry 110 V AC at all times. Accidentally touching the ends of the leads together might result in an unpleasant shock.

Uninterruptible Power Supply

An uninterruptible power supply (UPS) temporarily supplies power using batteries to the **load equipment** (anything connected to the UPS that draws power, usually servers and possibly other network equipment) in the event of a power outage. The UPS also supplements power in case of a brownout, where utility power continues but is below acceptable operating voltages.

A UPS typically provides backup power in that the load equipment constantly receives power from one or more backup batteries. The batteries receive a constant charge from utility power. If the utility power fails, the batteries continue to provide power just as they always have, minus the battery-charging function from utility power. The primary purpose of a UPS is not to continue normal operations for the entire duration of a brownout or power failure. Instead, the UPS provides a few minutes of power to give administrators enough time to send network messages to users (giving them time to save and close files) and gracefully shut down the server using normal procedures in the NOS. Otherwise, users can lose data from open files and the NOS can become corrupt or unstable.

In a Microsoft Windows environment, you can notify connected users individually or as a group. For example, at a command prompt you could enter "NET SEND /USERS Save Files Now and Log Off!" (Of course, sending users a message only applies if the power outage is limited to the server room.)

Power protection systems fall under three major architectures, as shown in Table 4-2.

Table 4-2 UPS Architectures

UPS Type	Function	Advantages	Disadvantages
Standby/Offline	A transistor momentarily switches a large transformer, which stores a small amount of power, before transferring power to the UPS battery. These UPS systems are better suited for home PCs, workstations, and so forth	• Low cost • Energy efficient—only converts DC current to AC during blackout or brownout	• No long-term brownout operation because battery power only lasts a few minutes • Provides minimal noise filtering on incoming power—does not regenerate power • Output power variance—can pass voltages outside normal levels up to +/-20%

Table 4-2 UPS Architectures (continued)

UPS Type	Function	Advantages	Disadvantages
Line Interactive	Similar to Standby/Offline, but adds automatic voltage regulation to stabilize power levels during brownouts and overvoltages without using battery power. Used in PC desktops as well as smaller servers, and usually provides up to 3 kilovolt/amps (kVA).	• Only nominally more expensive than Standby/Offline • Saves battery power by not operating during brownouts or overvoltages, instead using Automatic Voltage Regulation (AVR) circuits • Efficient energy use—only converts DC current to AC during blackout or brownout	• Provides minimal noise filtering on incoming power—does not regenerate power • Output power variance—can pass voltages outside normal levels up to +/–10%
Online/Double Conversion	AC utility power enters UPS where the rectifier converts to DC and then back to AC out to the load equipment. This conversion process cleans the power stream to near perfection. Used for higher-end, mission-critical servers, and is commonly used where more than 3 kVA is required*	• Precision output to within 2% of normal levels • No switchover time because load equipment already operates through the continuously charged battery • Extremely clean power output	• Reduced energy efficiency because the AC/DC/AC conversion reduces efficiency to about 85% • Expense—though costs are coming down, Online/Double Conversion UPS systems are still the most expensive

* Extremely large power requirements over 3000 VA often require a huge amount of space, and might be centrally located with specially wired UPS-supported building circuits or power distribution strips.

Several sources recommend that you do not purchase **standby power supply (SPS)** equipment, which detects an interruption in line power and switches to a transformer to bridge the period of time it takes to switch to battery power. This is no longer an issue, as it was in the 1980s and '90s, because the quality of UPS equipment and server power supplies can easily survive the momentary transition without ill effects. The switch might take 2–4 ms, while power supplies can usually handle a 100–200 ms pause.

Understanding UPS Operation

The following description and Figure 4-3 describe what takes place during normal operation of an online UPS when utility power is at normal, uninterrupted levels:

1. The UPS receives AC power.
2. The UPS uses a rectifier to convert AC power to DC power.
3. Some DC power is siphoned off to charge the battery.
4. An inverter converts DC power back to clean, nearly perfect AC power.
5. Power passes through the transformer to the load equipment.

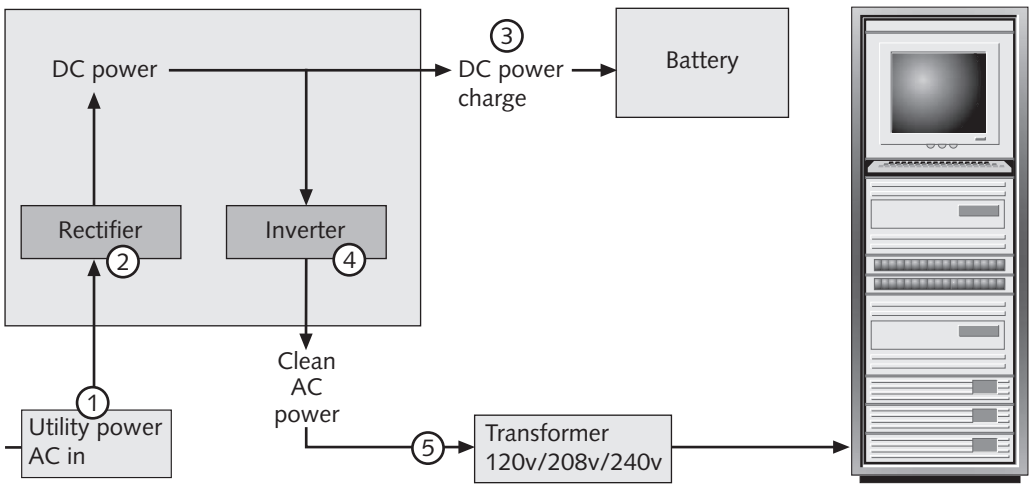

Figure 4-3 Normal online operation

The following steps and Figure 4-4 describe what takes place during a power failure:

1. Because there is no AC in, the battery discharges DC power to the inverter.
2. The inverter converts DC power to clean, nearly perfect AC power.
3. Power passes through the transformer to the load equipment.

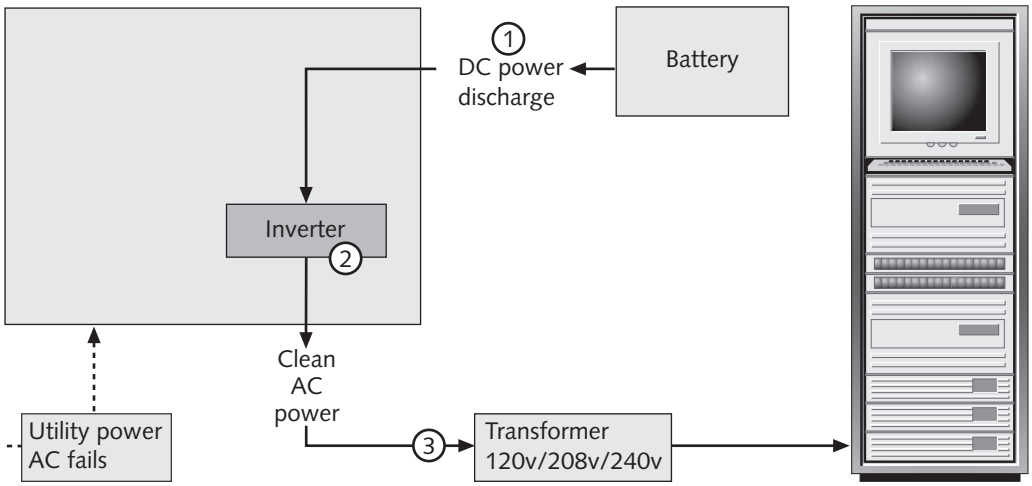

Figure 4-4 Battery operation

UPS Software

UPS software such as American Power Conversion (APC) PowerChute can also assist the administrator by automatically shutting down the server and safely storing data that might otherwise be lost. The UPS connects to a port on the server (usually a serial or USB port) and sends a message to the server when backup functionality activates. The message serves as a trigger to the software, which begins to perform administratively predetermined functions such as data backup and system shutdown. Many UPS manufacturers are including increasingly sophisticated administrative software that goes beyond these basic functions. For example, Tripp-Lite offers software that allows you to remotely manage UPS systems of most major manufacturers from an Internet connection.

Failover for the UPS

Backup batteries eventually fail or lose their ability to retain a charge. Server-level UPS systems usually also offer N+1 functionality that allows you to replace a battery while another battery (or batteries) continues to power the unit, ensuring that the UPS is not temporarily unavailable. Also, you can use multiple UPS units so that you can service one UPS while the remaining units continue to supply power (see Figure 4-5).

Generally, administrators seek about 15 minutes of backup power for servers. Depending upon the business need, administrators might seek up to eight hours of backup power, such as for PBX telecom systems and Internet connections. However, UPS equipment providing eight hours of backup is large and expensive. If you want backup power for an extended period of time, consider a backup generator with a UPS. The primary power service in the event of an outage comes from the generator, not the UPS. However, the generator usually takes several seconds to come fully online (known as **generator kick**, **kick**, or **kickstart**). The UPS in this context provides power during the seconds required for the generator to come online.

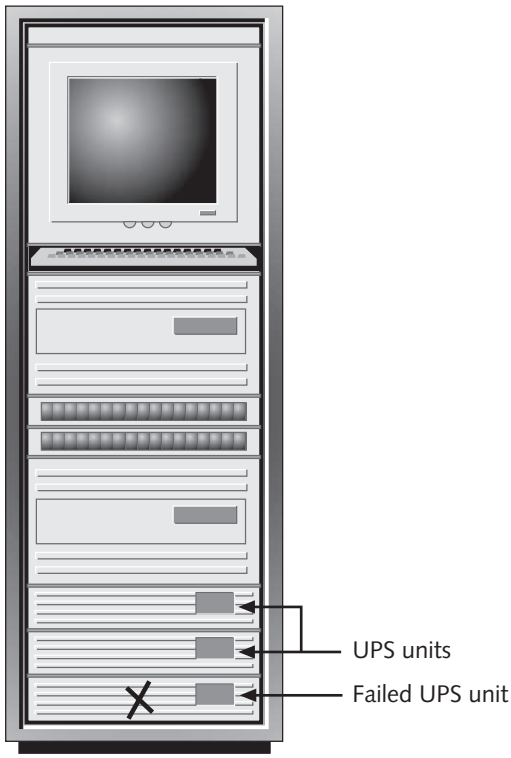

Figure 4-5 N+1 redundancy: if one UPS is unavailable, the remaining units continue to supply power

A **line conditioner** filters out power inconsistencies, temporarily bridges power in the event of a brief brownout, suppresses high voltage spikes, and provides overall buffering between utility power and the system. This chapter does not separately address line conditioners because most server-level UPS systems already include line conditioning capability.

Determining UPS Requirements

Determining your exact UPS requirements involves several variables, including power requirements of individual servers and networking equipment and/or rack power, one or more monitors, and so forth. The best way to calculate these needs is to add the total power requirements of the rack or server. Some sources recommend looking at the UL (Underwriters' Laboratory) sticker (if present) on the back of the unit to determine its total power requirement, but I do not. The UL sticker is a measure only of what the manufacturer submitted as the default configuration for this server, and servers often include additional equipment. Instead, measure the wattage required by your load equipment. Let's say that on a particular server, you need 600 W (which includes 50% overhead, as recommended earlier). When you visit the web site of a UPS manufacturer, you

only see UPS systems rated in volt-amps (VA), and notice that a volt-amp is calculated using the same formula as watts: volts × amps. Since the formula is the same, does that mean watts and VA are the same thing? No, they are not. The difference is "where" the power occurs.

Power coming "in" from the utility company is measured in watts. That's what you pay for on your electric bill, and it is also sometimes called **actual power** or **true power**. However, as the electricity passes through the server's power supply, capacitors, inductors, and other equipment, we must account for a difference between the power that comes "in" to the power supply and power that goes "out" of the power supply. Power going "out" is known as **apparent power**. The difference between actual power and apparent power is known as the **power factor**, which is usually a difference of about 60%.

To boil it all down, it comes to a simple factor in determining UPS power requirements. Take the VA rating of the UPS and multiply that number by .60 to determine the number of watts that this UPS will support. For example, if the UPS is 1000 VA, then the watt rating for the same UPS is 600 W (1000 × .60 = 600 W). In the example stated at the beginning of this discussion, a 1000 VA UPS would exactly meet the 600 W power requirements of the server equipment. However, to plan for future expansion and add a margin of safety, you should increase to the next higher available VA rating from the UPS manufacturer.

Next, determine the amount of time (known as the **run time**) you require to power the server. Realize that VA × high run time = lots of money. Run time is not a calculation you make on your own; you have to contact the UPS manufacturer to make that determination. This is best accomplished by visiting a UPS web site and using online tools to arrive at the best product for your needs. (Try Hands-on Project 4-4 at the end of this chapter.)

Avoid including printers in the total UPS power requirements, because documents can usually be printed any time; risk of data loss is a more immediate concern. Also, printers can be electrically "noisy," drawing varying levels of power; when performing a print operation, they are extremely demanding, electrically speaking. (Inkjet printers require much less power than laser printers and are better to include on a UPS if you absolutely must have a printer during a power outage.)

Site Preparation

In addition to calculating the total power requirements for your servers, you must also consider the physical space that a UPS requires. Batteries can be quite large, and several models allow you to daisy chain multiple external battery packs together to extend UPS run time. This can require a great deal of space. One solution is a rack-mounted UPS. All vendors are moving to rack-mounted models in various formats. You can connect UPS units either in the same rack or in adjacent or back-to-back racks to increase run time or provide UPS fault tolerance (see Figure 4-6a). Also, you can use a rack-mounted

UPS with an external battery pack (see Figure 4-6b). These solutions are not something administrators arbitrarily piece together. Instead, always consult the UPS manufacturer, which will advise you on permissible physical connections, safety, power capacity, and run time.

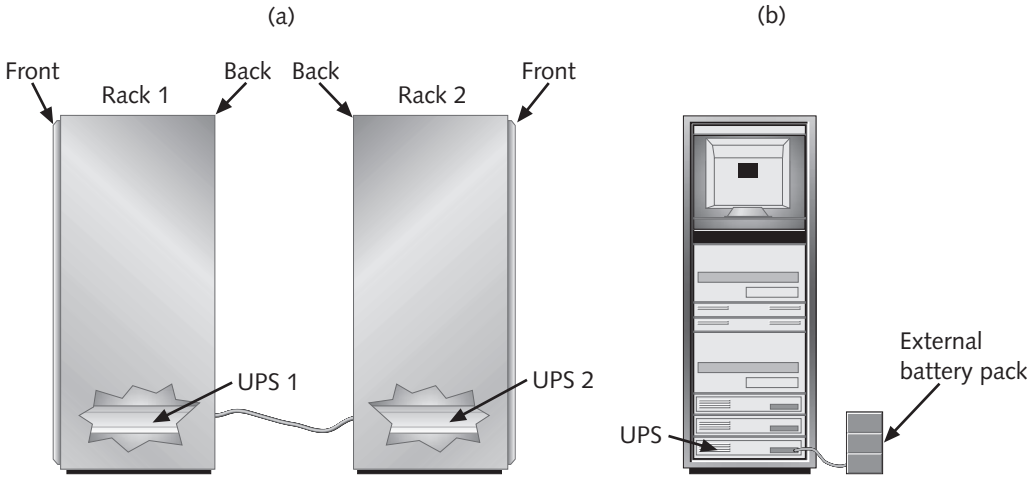

Figure 4-6 To provide extra power, connect two UPS systems back to back (a) or use an external battery pack (b)

It is important to realize that adding batteries to the UPS does not increase its volt-amp rating. For example, adding an extra battery will not upgrade your 1000 VA UPS to 1500 VA. Instead, it only extends run time. If the run time is 15 minutes, adding an extra battery will power the 1000 VA UPS for a few more minutes. The only way to increase VA is to use a more powerful UPS.

Rack-mounted UPS systems, especially larger ones, are extremely heavy and usually require assistance to install and service. Because of the weight, you should place them at or near the bottom of the rack (see more about rack-weight distribution later in this chapter).

The largest UPS systems can be huge. For example, the APC Silicon DP3500E is 70 inches high, 94 inches wide, 31 inches deep, and weighs a staggering 5500 pounds. Regardless of the size, the following checklist will help you to plan site requirements, particularly for larger systems:

- For larger UPS systems, notify building engineers, electricians, and electrical engineers; for extremely heavy systems, also notify structural engineers. Verify that all aspects of the site are safe.
- Do not place monitors or other devices that are highly sensitive to electromagnetic fields (EMF) near the UPS. A larger UPS might emit an EMF that affects computer monitors, but probably not other equipment.

- UPS systems generate heat. Make sure the HVAC system can accommodate the heat output.
- Plan for any rewiring so that load equipment can reach the UPS. To prevent overloading the circuit, place larger UPS systems on a single, dedicated circuit.
- Verify sufficient clearance for adequate airflow, probably around 12 inches (0.31 m) behind a freestanding UPS array and 48 inches (1.22 m) in front (see Figure 4-7). Ensure adequate space to reach all switches, jacks, outlets, and so forth.

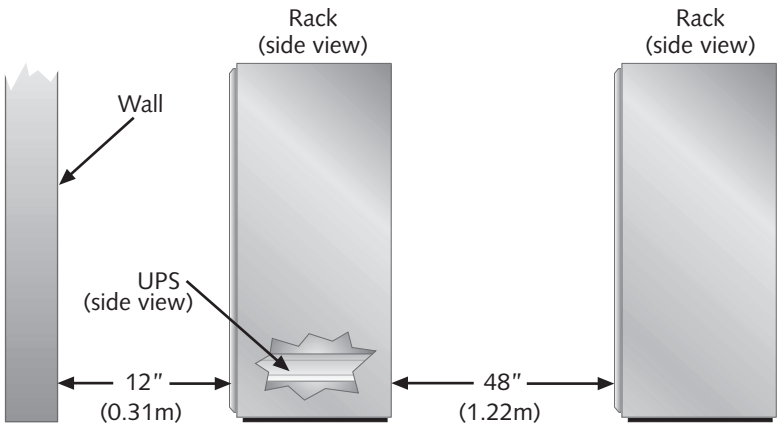

Figure 4-7 Allow enough space for airflow and equipment access

- Many UPS units have an optional earth ground to which you can ground the unit (in addition to existing AC grounding). Locate a suitable ground location reachable from the UPS.
- Get help—even in smaller units, the batteries are surprisingly heavy.
- Some larger freestanding units are on rolling casters to make movement easier. Make sure you lock the wheels when not moving the unit.

The following is a list of major vendors of UPS equipment:

- Tripp-Lite—*www.tripplite.com*
- American Power Conversion (APC)—*www.apc.com*
- Best Power—*www.bestpower.com*
- Liebert—*www.liebert.com*
- MGE—*www.mgeups.com*
- Oneac—*www.oneac.com*
- Sutton Designs—*www.suttondesigns.com*

THE RACK

In addition to configuring the system itself, you must determine the best physical orientation for the server—freestanding or rack mounted. Except for the smallest networks, you should plan to install servers in racks (this is the general assumption of this and following chapters unless stated otherwise). Racks provide an advantage in server rooms where the need for server and network equipment grows but floor space does not. By stacking equipment in a rack, you increase computing assets vertically in the same floor space that would have otherwise consumed precious floor space. This space savings is known as **density**. For example, in Figure 4-8, six servers on a table require about six feet of floor space. Using racks in the same six feet, you can install 18 8-way servers.

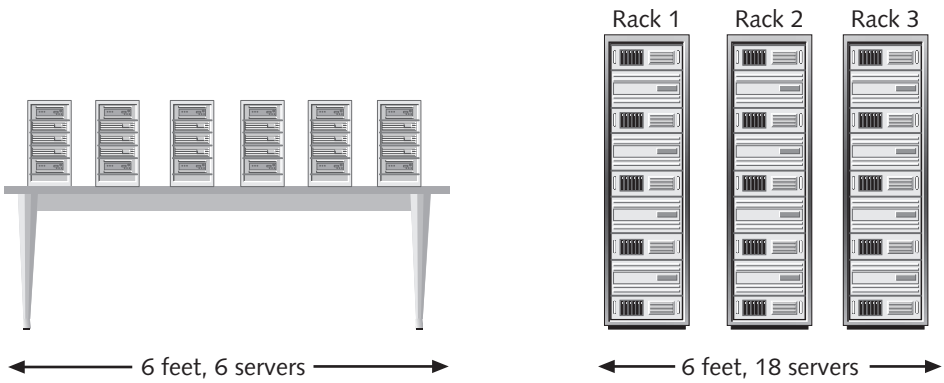

Figure 4-8 Density allows for more equipment in the same horizontal floor space

If you have existing servers in the tower configuration and want to place them in a rack, you can either place them on a vented shelf that mounts inside the rack or obtain a vendor kit that adapts the server to lay on its side and install it in the rack along telescoping rails. Nearly all network equipment (except, perhaps, for small workgroup hubs or switches) is rack-mountable. When you decide to use a rack, you must carefully consider several factors, including heat, ventilation, power, weight, grouping, and accessibility.

Physical Characteristics

Physical rack characteristics vary from one vendor to the next. However, general characteristics are listed here. All aspects of equipment installation revolve around the physical dimensions of the rack as follows (see Figure 4-9):

- *Units*: Rack equipment is measured in **EIA (Electronic Industries Alliance) units**, or **U**. One EIA unit (1U) is equal to 1.75 vertical inches (4.45 cm). For example, Tripp-Lite makes a rack-mounted UPS that is only 1U in size. Server appliances are usually between 1U for network appliances

dedicated to a single purpose and up to 7U (12.25 inches, or 31.12 cm) for 8-way servers. Datacenter servers larger than 8-way (such as a 32-way) are usually in their own dedicated enclosure.

- *Height*: Including frame, bezels, and feet or rolling casters, a full-height rack (42U) is about 6 feet (1.8 m). Various manufacturers make smaller racks as well, with common sizes at 22U, 24U, and 36U. Some manufacturers offer rack extensions, which add about 8U to the height. See Table 4-3 as a reference (casters or feet not included).

Table 4-3 Rack Heights in Units, Inches, and Centimeters

Height in Units	Height in Inches	Height in Centimeters
22U	38.5	97.8
24U	42.0	106.7
36U	63.0	160.0
42U	73.5	186.7

- *Depth*: Racks are about 36 inches (.98 m) deep, with usable depth around 28 or 29 inches (between .71 and .73 m). Try not to use space beyond the usable depth, because you will still need room at the back to work, and for PDUs, cables, and other specialized devices that you do not need to see from the front.

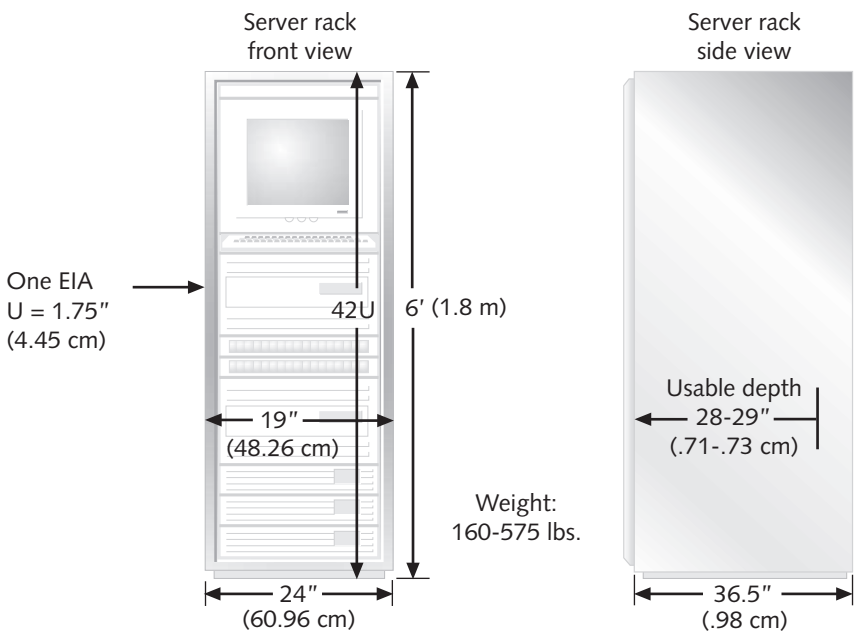

Figure 4-9 42U rack dimensions

- *Width*: Rack-mountable equipment requires a RETMA (Radio Electronics Television Manufacturers Association) industry standard opening of 19 inches (48.26 cm). Exterior width is usually about 23–24 inches (58.5 cm).
- *Weight*: Weight varies depending upon whether the rack includes ballast or doors, which weigh about 30 pounds each. 42U racks can weigh around 160–250 pounds empty, and can accommodate about 2000 pounds of equipment. (IBM has a heavy-duty rack that weighs 575 pounds empty!) 24U racks are about 210 pounds empty, and can accommodate about 1000 pounds of equipment.

Racks can also include several other items:

- *Stabilizing feet*: Moving rack equipment in and out while servicing it presents a physical danger of the rack falling over. Stabilizing feet extend beyond the rack and help prevent it from tipping over when you work on equipment. The HP Rack System/E includes both a front and rear retractable anti-tip foot for added safety.
- *Leveling screws*: Similar to the leveling screws under a washing machine (but capable of much greater weight loads), you turn the screws until the rack is level.
- *Wheels*: Usually made of polyurethane to reduce the effect of bumps, jolts, and uneven floors when moving the rack. The wheels can support a great deal of weight, usually 1000 pounds *each*.
- *Filler panels*: Purely cosmetic, these panels cover up empty slots for a more professional, finished look.
- *Side panels*: Also cosmetic, side panels cover up the exposed sides of the rack. These are unnecessary where you join two racks together.
- *KVM/concentrator/switchbox*: A **keyboard, video, mouse (KVM)**, also called a concentrator or switchbox, enables you to control multiple servers from a single keyboard, video monitor, and mouse.
- *Cable management arm*: It is important that when you service equipment and pull it out of the rack, all the cables do not come loose. Instead of bunching up the cable slack in a tangled mess at the back, you can Velcro cables to a **cable management arm (CMA)** that keeps cables neat while allowing them to extend when you pull out equipment. Cable management brackets also can be used to guide cables vertically within the rack.
- *Ballast*: Dead weight placed vertically, at the sides of the rack or at the bottom of the rack, to add stability when heavy equipment is required higher in the rack. A single ballast usually weighs about 30 pounds.
- *Short rear door*: The rear door has a gap of a few inches at the bottom to facilitate cabling out of the rack while maintaining security.

Cooling

Many racks include an option for front and rear lockable doors. Most doors are perforated safety glass or steel, providing about 60% opening for adequate ventilation, which in most racks is through convection (see Figure 4-10). Some manufacturers (such as Compaq) offer a multi-angled door design that enhances convection cooling. Warm air rises to the top of the rack, which may also be perforated.

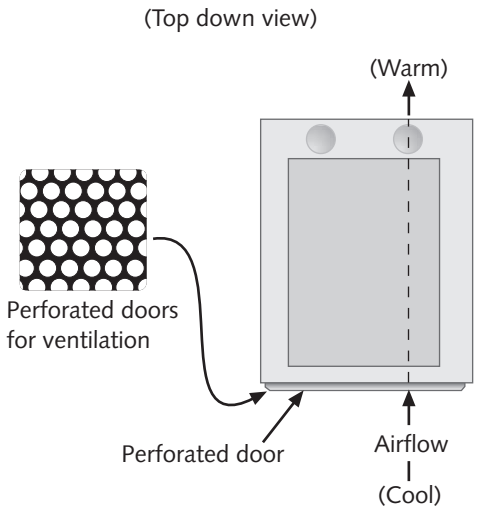

Figure 4-10 Perforated doors assist cooling through convection and fan-propelled airflow through equipment

To ensure adequate cooling in the rack, you can install fans in the top panel to increase availability of cool air inside the rack and draw warm air out (see Figure 4-11). Other temperature-control solutions cool from the front to the back, which is a function of the equipment in the rack. For example, a server might have a fan at the back of the unit that draws air into louvers at the front, and expels warm air out the back. Even in this case, you might still install fans at the top, which will help to more quickly expel naturally rising warm air while also drawing cooler air from the bottom of the rack. (Recall from Chapter 2 that some floors have air-conditioning vents beneath the racks.) This is also a more significant concern when the rack includes items that run hot, such as UPS systems and disk arrays.

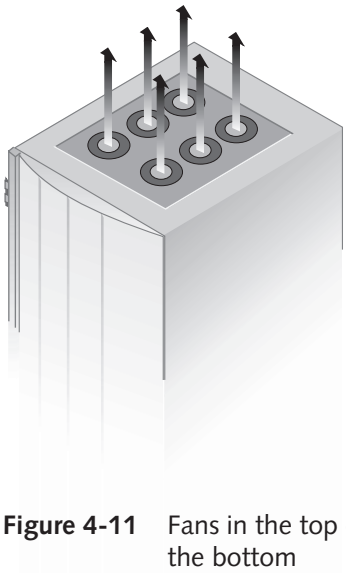

Figure 4-11 Fans in the top of the rack draw warm air out while cool air enters from the bottom

Be careful not to place anything on the top surface of the rack, which would inhibit dissipation of warm air. Be careful not to place racks so that the back of a rack faces the front of another to avoid the intake of warm air. Instead, you should place racks back to back and allow adequate space between rows of racks (see Figure 4-12).

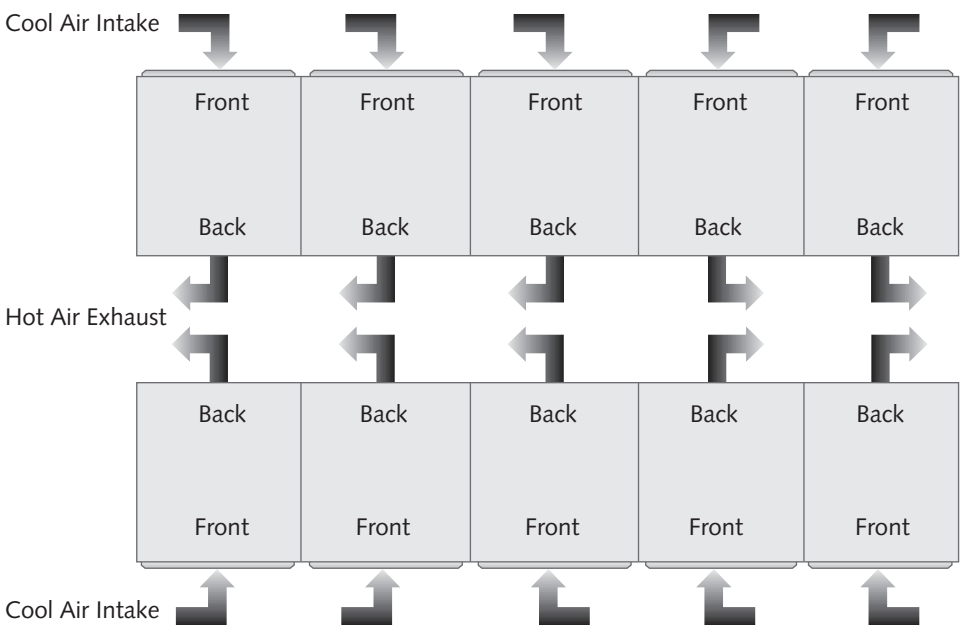

Figure 4-12 Arrange rows of racks in a back-to-back configuration

Although a rack with doors can also be classified as a cabinet, the remainder of this book classifies both types as a rack unless specified otherwise.

Configuration

Configuring the rack well directly affects the level of practical usability and prevents the need to rearrange equipment in the future. Major vendor web sites offer rack configuration utilities as either downloadable programs or web-based applets. These utilities are the most effective way to configure a rack, especially if all rack equipment is from the same vendor. If you make an unwise choice, the utility might alert you and ask if you want to select an alternative. Even if you do not stay with the same vendor for all equipment in your rack, you can use a rack configuration utility to approximate like equipment. For example, the HP Netserver LXR8500 is very similar in power requirements and dimensions to the Dell PowerEdge 8450. The following is a short list of vendors offering rack configuration utilities:

- Dell—*www.dell.com* (download). A good, basic rack configuration utility.
- Compaq—*www.compaq.com* (web-based). Offers the advantage of not having to install yet another program on your computer. Offers two modes, one for novices and one for those more experienced in rack configuration.
- Hewlett-Packard—*www.hp.com* (download). This one is excellent in terms of usability and has a low learning curve. In Hands-on Project 4-5 at the end of this chapter, you will use the HP Rack Assistant utility to practice configuring a rack.
- IBM—*www.ibm.com* (download). The most detailed tool, but it will take a little longer to learn its usage and options (see Figure 4-13).

While extremely useful and educational, these utilities cannot account for every contingency, and you will have to monitor the results for accuracy and practicality in your own real-world environment. Also, most of these utilities do not fully consider the implications of joining two racks together, especially in terms of sharing cabling between the racks. Therefore, you should be able to wisely configure a rack for weight distribution, device grouping, and cable management as described in the following sections.

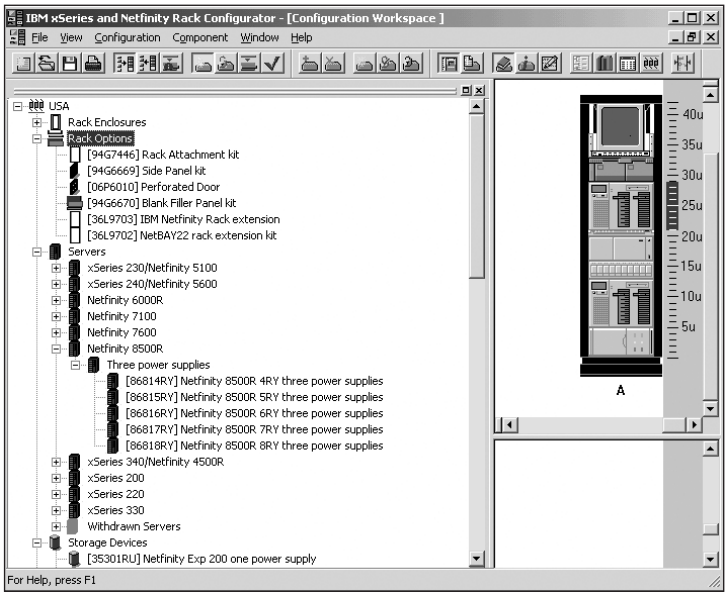

Figure 4-13 The IBM Netfinity Rack Configurator is an excellent tool

Weight Distribution

Larger 8-way rack servers can weigh upward of 175 pounds, and UPS systems like the APC Smart-UPS 5000 RM is 5U and weighs 320 pounds! The general rule of thumb is to *place the heaviest items at the bottom of the rack*, which minimizes the chances of the rack tipping over. Therefore, unless the UPS is small (1U or 2U), you should almost invariably place it at the bottom. If other planning factors such as device grouping prevent you from placing all the heavier items at the bottom, try to avoid placing heavier items any higher than 36 inches (.91 m) up the rack, and consider adding one or more ballasts. Each ballast is 1U and weighs 30 pounds. Other heavy items include mass storage items such as a DLT (Digital Linear Tape, a backup device) and a disk array. (HP recommends placing the DLT just above the UPS.) It should not adversely affect weight distribution if you need to insert a keyboard/mouse between heavy devices.

It is very important that you pull out only one piece of rack equipment at a time to avoid tipping. Even equipment that is only moderately heavy has a greater impact on server balance when it is pulled "out of center." A rack can easily exceed 1000 pounds, and could seriously injure or kill someone if it fell over. When working on a rack, don't forget to extend stabilizing feet if they are available (see Figure 4-14).

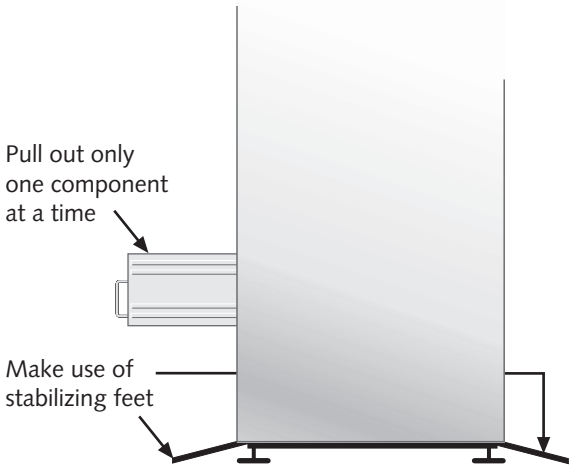

Figure 4-14 Pull out only one device at a time, and use stabilizing feet to prevent tipping

Device Grouping

Grouping devices involves a delicate balance between weight distribution and logically placing components where cables from one component can reach other components. Also, for purposes of usability, some types of equipment should be placed where they are most usable. For example, although the keyboard is lightweight, placing it at the top of the rack would make it inaccessible to administrators. If administrators work on the server in a standing position, you should place the keyboard in the middle of the rack (or a little lower if administrators will be seated). The logic of some devices is also important. For example, certain devices might need to be closer to one another for cables to reach.

Another reason to group equipment in a different way than you would if only considering weight might be clustering. As introduced in Chapter 1, clustering involves two or more servers serving data from the same physical media (hard drives). In the rack, you might have (starting from the bottom) at least one UPS, Cluster Server 1, a disk array that contains data, and then Cluster Server 2 (see Figure 4-15).

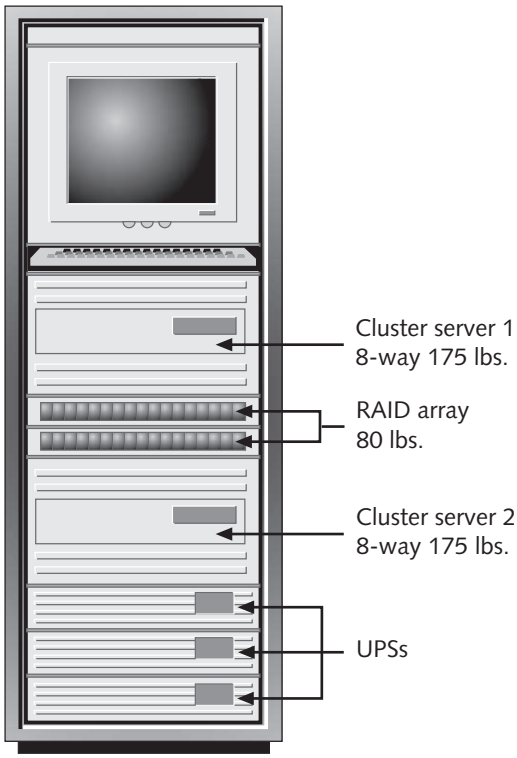

Figure 4-15 A cluster might change the usual weight distribution order

Cable Management

One of the primary goals in cable management is neatness. While this might sound compulsive, organization is critical in the server room. When equipment fails, administrators must have minimal distraction in finding the proper equipment. When trying to access equipment at the back of the rack (to replace a power supply, for example), you do not want to fight your way through a morass of cable. With cables neatly aligned along a CMA and placed out of the way with cable management brackets, guides, and so forth, you should be able to quickly access the equipment you need. Figure 4-16 is a good example of neatly arranged cable.

Another factor that affects where you place equipment is cable length. The effective cable reach of pieces using a CMA will be shorter because several inches (perhaps 28 or more) will be used by the arm. However, the shorter cable reach is well worth the improved organization and secure connections that a CMA provides (see Figure 4-17).

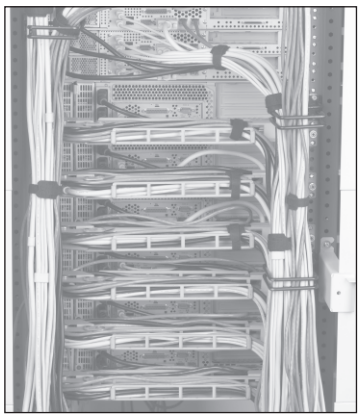

Figure 4-16 An HP rack with neatly arranged cable

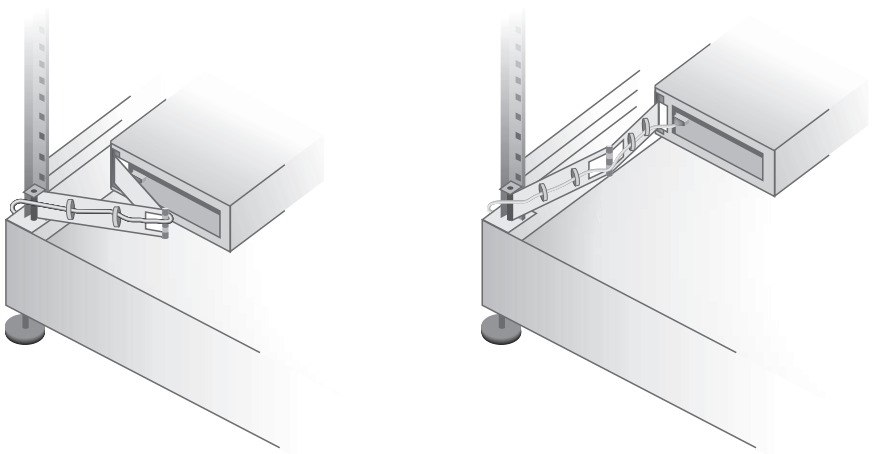

Figure 4-17 The CMA expands and contracts when you move the server

If cable lengths are too short, you might have to place dependent equipment closer to one another. For example, a light, 1U web server near the top of the rack must back up its content to a DLT located at the bottom of the rack. If the cable lengths are not sufficient, you might have to move the heavy DLT higher in the rack, lower the 1U web server, or move the 1U web server to an adjacent rack in a lower position closer to the DLT. Many devices cannot use an extension to make up the shortfall in cable lengths. For example, there is no such thing as a SCSI cable extension.

Cable lengths become a more visible restriction when you consider a monitor, which is usually placed higher in the rack than the server. Typically, the monitor cable can reach

the server over a maximum of 29U. If necessary, you can purchase an inexpensive VGA extension cable. I recommend you use as short a length as necessary (hopefully 6 feet or less), because longer lengths are not good for video signal integrity.

In the absence of a CMA, you might consider using a straight point-to-point cabling method with no intermediate cable management. However, this is the least desired method because the rack becomes more difficult to manage when you have cables hanging in the way. Also, unless there is a method to secure the cable ends (such as thumb screws), they are more likely to fall off. Some rack equipment has a cable tray option extending from the back of the equipment. The point of the cable tray is to lessen the pull of gravity on the cable connection and reduce the likelihood of the cable falling out (see Figure 4-18).

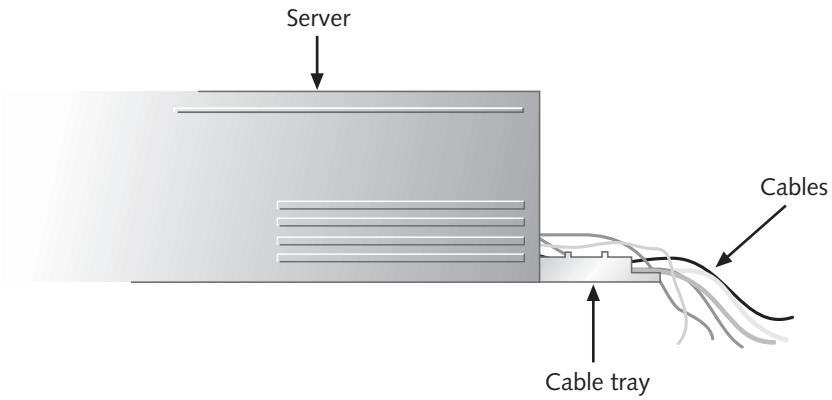

Figure 4-18 A cable tray helps cables to stay plugged in

Sometimes you need to use multiple devices from several racks that connect to a physical and logical center. For example, Figure 4-19 shows a centrally located server that connects to several devices such as DLT backup equipment, a SCSI disk array, and a Fibre Channel disk array. If other planning factors prevent you from placing all the equipment in the same rack, then centrally locate the server between racks containing the equipment.

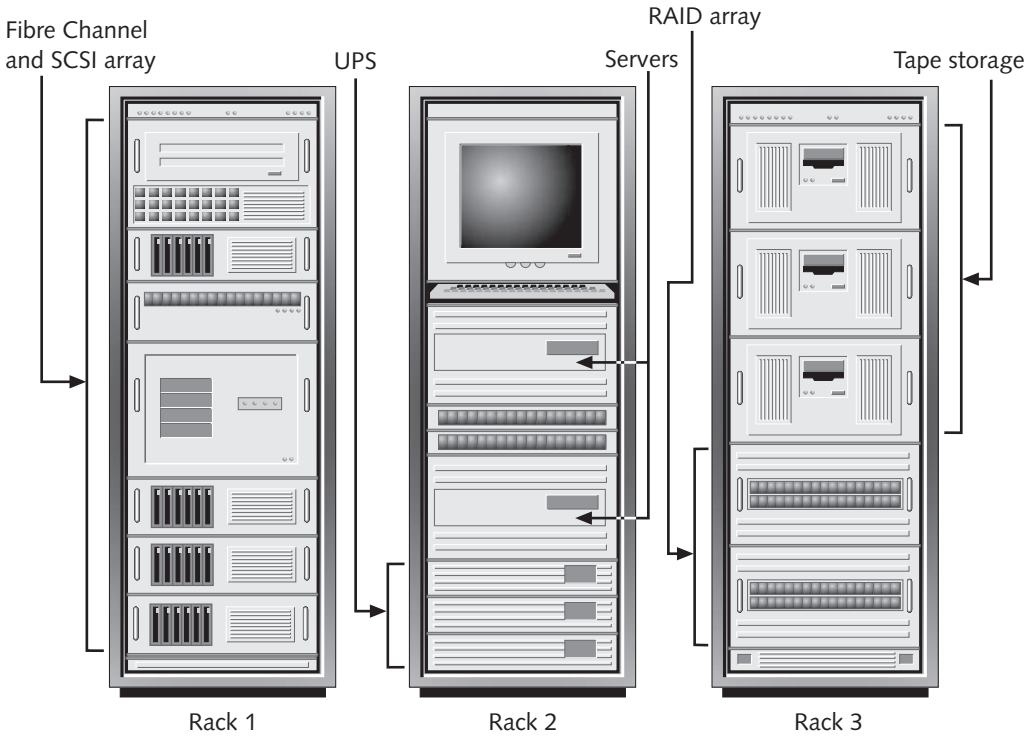

Figure 4-19 Place the server central to the other equipment to which it connects

Be sure to route cables so that they cannot be stepped on or tripped over, and do not place any equipment on top of the cables.

Power

The need for a UPS was established earlier in this chapter. A UPS is only as good as its availability, so each piece of rack equipment must be able to connect to the UPS. The UPS has a limited number of outlets (perhaps six), and in a fully loaded rack, that's probably not enough. However, by using a rack-mounted PDU, you can create more outlets (see Figure 4-20). Make sure that the devices you attach to it do not exceed 80% of the PDU's power rating. You install the PDU in the rack in any of several configurations. Usually, you try to locate the PDU at the rear of the rack and as close to the bottom as possible.

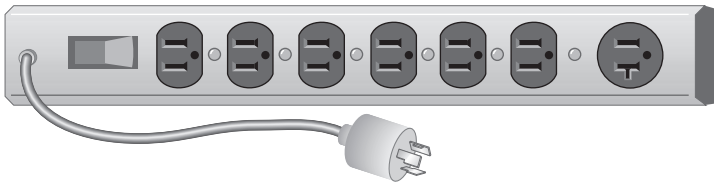

Figure 4-20 The PDU provides the functionality of a power strip at much higher power capacity

Vertically installed PDUs attach to either the left or right rack post (see Figure 4-21). However, most PDUs use 11U vertically, so you can probably only install three on one rail, left or right. Another factor that limits PDU installation is the type of server. Some servers slide out of the rack at the rear instead of the front (the HP NetServer LXr Pro8, for example). In this case, a horizontal PDU in the same space as the server would interfere with server removal. In a vertical PDU configuration, also be careful not to install the PDU at the location of the locking latch, and face the PDU inward so that rack covers do not block outlets.

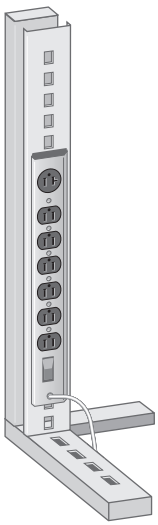

Figure 4-21 A vertical PDU installation

A horizontal installation usually allows for cleaner routing of power cables, and can be mounted behind most rack equipment in the same EIA unit except for very deep, large rack components such as some large servers. If installation behind a unit is not acceptable, you can also mount PDUs in the bottom of most racks (see Figure 4-22).

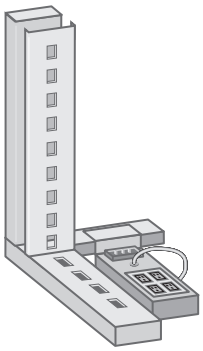

Figure 4-22 Mounting in the bottom of the rack does not interfere with other equipment

Many PDUs, especially those with higher voltage, have locking plugs that plug into a locking receptacle to prevent accidental disconnection. The rack utility software from several vendors mentioned earlier calculates the VAs required to service the equipment you propose, and also suggests the quantity and type of PDUs you will need. For an added level of power redundancy, you can use a PDU designed for utility power from two independent circuits. If one of the circuits fails, the other transparently continues to provide service (see Figure 4-23).

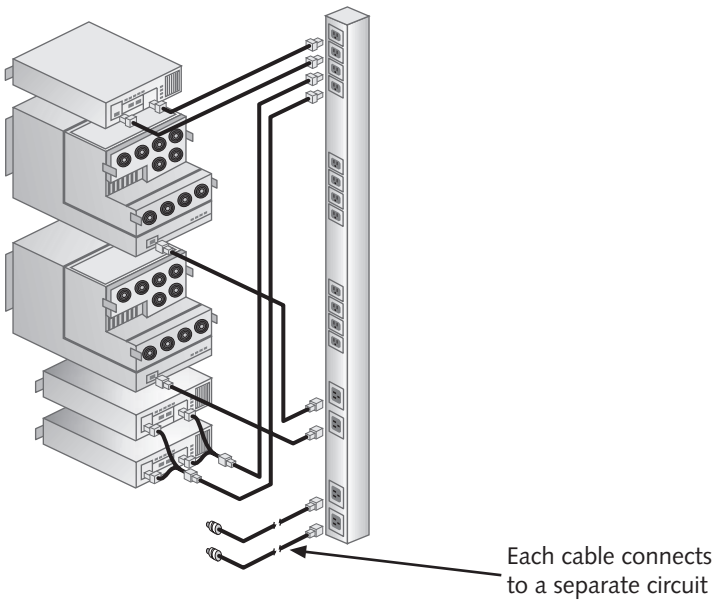

Figure 4-23 For added redundancy, use a PDU that connects to two independent circuits

Turn on components in the rack one at a time to avoid a sudden inrush of power, which might trip the main circuit breaker, UPS fuse or breaker, or PDU fuse or breaker. Usually, you should power up in this order: tape backup, mass storage units, monitor, KVM, and one server at a time.

The KVM

It is impractical to have a separate keyboard, video display, and mouse for each server in the rack. A 42U rack can have 42 separate 1U servers! Instead, it makes more sense to install a single keyboard, video display, and mouse (KVM) that can service all of the connected servers, collectively referred to as a **console** (see Figure 4-24). A KVM console also reduces air-conditioning costs by eliminating multiple heat-generating monitors. You can obtain a basic, inexpensive KVM console for any configuration of workstations or servers, most of which operate between four and eight systems. Some can also be set to continually cycle between servers every few seconds so that you can observe activity on each server.

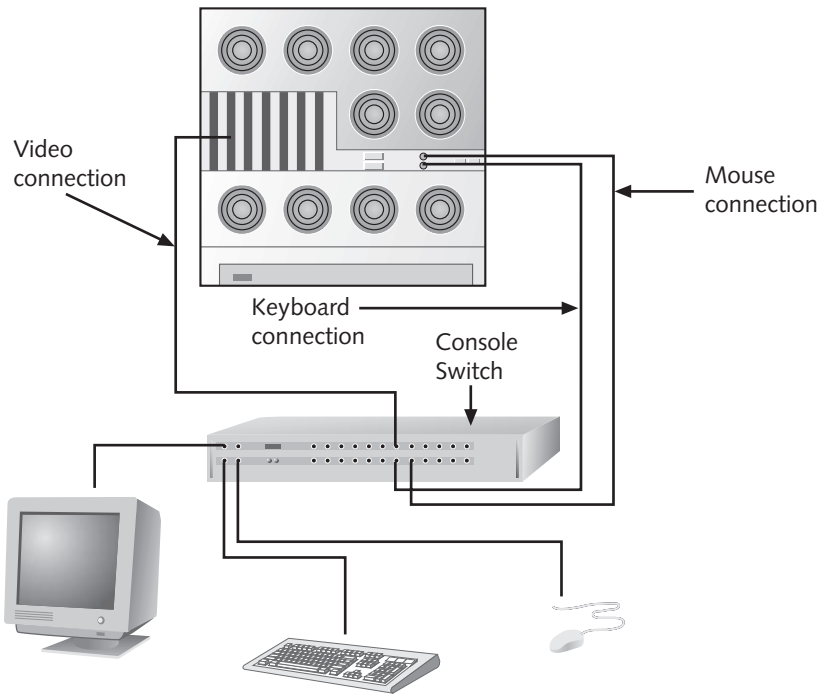

Figure 4-24 Use one keyboard, video, and mouse to control servers through a KVM console

For the rack, you can use several KVM components to arrive at the configuration you want:

- 1U fixed keyboard/mouse tray
- 2U retractable keyboard and mouse tray
- Full-size (about 11U) video display
- 1U or 2U integrated keyboard, trackball, and video display (see Figure 4-25)

Figure 4-25 An integrated keyboard, trackball, and pop-up video display

Whichever configuration you decide on, the clear advantage is that you conserve space by using a single KVM console to manage multiple servers. Also, since each KVM console often manages up to eight servers, vendors usually offer the option to cascade the consoles, which allows you to manage dozens of servers. Some consoles are wireless, and with signal amplification can allow you to control more than 60 servers.

The keyboard, mouse, and video monitor each connect to the KVM console. Then you need one extension cord (male to female) each for the keyboard, video monitor, and mouse for each server. If you have four servers, then you need four sets of extension cords. You can usually purchase the set as a bundled pack from any computer supply source.

KVM switches are available from nearly all major server vendors and most computer stores. APC manufactures an outstanding KVM console that, in addition to standard functions, includes the following features:

- Password security to the KVM console adds an extra level of security.
- An on-screen display with customizable menus allows you to graphically switch between servers.
- Hot-pluggable operation allows you to add servers without having to first power off the KVM or other servers attached to the KVM.
- Mouse reset circumvents a frustration with KVM systems in which the mouse ceases to respond to actions. This feature also allows you to regain control of the mouse without powering down the KVM or the server.

Alternatively, you can usually administer servers from the comfort of your own desktop with a full-screen monitor, keyboard, and mouse, depending on the operating system. For example, you can use the text-based Telnet utility to remotely connect to and administer

UNIX systems and perform router programming. Windows 2000 and NT 4.0 offer administrative tools that you can run from any Windows computer on the network. NetWare offers similar management control.

Rack Installation Tips

The following tips can help you to properly configure the rack:

- Begin building the rack from the bottom up, with heaviest devices at the bottom (as previously discussed).
- If using multiple racks, be consistent in the server numbering scheme in relation to port numbers on the KVM, which helps you to quickly access the right server from one rack to the next instead of using a hit-or-miss method.
- Place servers attached to the same KVM console in close proximity to it so that you can easily view server activity while at the keyboard.
- Group servers that serve a particular purpose in the same physical area. For example, put web servers in one location, file servers in another, application servers in another, and so forth. This helps to minimize trips across the server room when configuring or troubleshooting.
- Remember that at times you must access the rear of the server. Make sure there is adequate space behind the server to open the rear door.
- Place servers that require frequent access (such as those with backup devices) in more easily accessible locations.
- Place servers according to security need. Position high-security servers in highly visible locations where it is more obvious if someone is accessing them.
- For added security, lock the rack doors and install an alarm that trips when the door has been breached.
- Fully loaded racks are extremely heavy, so limit the number of racks that you tie together if you plan to move them, even for routine cleaning. It can be very difficult to roll three attached racks, even with assistance.
- When sliding equipment in and out of a rack, watch your fingers! The slide rails are pinch points.

Hundreds of pages about rack planning and installation can be found at the following sites:
http://netserver.hp.com/netserver/docs/download.asp?file=g_rack_cabling.pdf
http://support.dell.com/docs/systems/smarcon/en/index.htm
ftp://ftp.compaq.com/pub/products/storageworks/techdoc/racksandoptions/14255e2.pdf
www.ibm.com (search for the Netfinity Rack Configurator)

CHAPTER SUMMARY

- A server can usually operate acceptably with only a single PSU. A second PSU provides load balancing to reduce demands placed on a single PSU and also provides failover, in case one of the power supplies fails.
- Servers offer N+1 expandability for several critical components, particularly the processor and the PSU. The "N" is a variable that refers to the quantity of a given component installed in a system, and "+1" refers to a spare component or opening for a component.
- The PSU often uses a variable fan speed, which increases or decreases based upon the amount of heat detected by the PSU thermistor (a type of internal thermostat).
- Power Distribution Units (PDUs) often plug into a 208 V outlet and supply power to internal rack components.
- Calculating the power needs for a server requires you to know how much power the motherboard, processor, internal adapters, and peripherals require.
- Some devices might list volts, watts, amps, or a combination thereof. Use the information you can gather in the following formula: watts = volts × amps.
- Add the total power requirements for all server components, and subtract the total number from the power rating for the power supply. You should have plenty of power to spare if you want the system to be as reliable as possible.
- An uninterruptible power supply (UPS) temporarily supplies power using batteries to the load equipment (anything connected to the UPS that draws power, usually servers and possibly other network equipment) in the event of a power outage. The UPS also supplements power in case of a brownout, where utility power continues but is below acceptable operating voltages.
- The primary purpose of a UPS is not to continue normal operations for the entire duration of a brownout or power failure. Instead, the UPS provides a few minutes of power to give administrators enough time to send network messages to users (giving them time to save and close files) and gracefully shut down the server using normal procedures in the NOS.
- UPS software can assist the administrator by automatically shutting down the server, issuing alerts, and safely storing data that might otherwise be lost.
- Backup batteries eventually fail or lose their ability to retain a charge. Server-level UPS systems usually also offer N+1 functionality that allows you to replace a battery while another battery (or batteries) continues to power the unit, ensuring that the UPS is not temporarily unavailable. Also, you can use multiple UPS units so that you can service one UPS while the remaining units continue to supply power.
- Generally, administrators seek about 15 minutes of backup power for servers.

- Determine the UPS requirements based upon the power supply rating and the amount of time you require to power the server.
- The following formula determines VA: volts × amps = VA, and 1000 VA = kVA.
- Rack-mounted UPS systems, especially larger ones, are extremely heavy and usually require assistance to install and service. Because of the weight load, you should place them at or near the bottom of the rack.
- By stacking equipment in a rack, you increase computing assets vertically in the same floor space.
- Rack equipment is measured in EIA (Electronic Industries Alliance) units (U). One U equals 1.75 inches (4.45 cm).
- The standard rack is 42U high and has a 19 inch (48.26 cm) opening for equipment.
- Stabilizing feet extend beyond the rack and help prevent the rack from tipping over when you work on equipment.
- Many racks include an option for front and rear lockable doors. Most doors are perforated safety glass or steel, and provide an opening for adequate ventilation, which in most racks is through convection. You can also install fans in the top.
- The general rule of thumb with rack installations is to *place the heaviest items at the bottom of the rack*, which minimizes the chance of the rack tipping over. It is very important that you only pull out one piece of rack equipment at a time to avoid tipping.
- Clustering, cable reach, and usability might force exceptions to the rule that the heaviest items always go on the bottom of the rack.
- Keep cables neat and within adequate reach of the devices to which they connect. Use a cable management arm (CMA) to prevent equipment from being unplugged when sliding them out of the rack.
- Make sure that the devices you attach to a PDV do not exceed 80% of the PDU's power ratings.
- Vertical PDU installations allow installation of other equipment in the same EIA units. A horizontal installation usually allows for cleaner routing of power cables and can be mounted behind most rack equipment in the same EIA unit, except for very deep, larger rack components such as some larger servers. You can also mount PDUs in the bottom of most racks.
- For an added level of power redundancy, you can use a PDU designed for utility power from two independent circuits. If one of the circuits fails, the other transparently continues to provide service.
- Turn on components in the rack one at a time to avoid a sudden inrush of power, which might trip the main circuit breaker, UPS fuse or breaker, or PDU fuse or breaker. Usually, you should power up in this order: tape backup, mass storage units, monitor, KVM, and then one server at a time.

- It makes sense to install a single keyboard, video, mouse (KVM) console that can switch back and forth between the connected servers. A KVM console also reduces air-conditioning costs by eliminating multiple heat-generating monitors. Vendors usually offer the option to cascade the consoles, allowing you to manage dozens of servers.
- You need one extension cord (male to female) each for the keyboard, video monitor, and mouse for each server.

KEY TERMS

actual power (also **true power**) — The power in watts delivered from the utility company.

apparent power — The power delivered to a device after passing through the power supply.

cable management arm (CMA) — Rack equipment that allows orderly arrangement of cables, and expands and contracts so that you can move equipment on the rack without accidentally unplugging it.

console — An inclusive term for the keyboard video mouse (KVM) and all attached servers.

density — A measure of the number of devices or servers within a given area of floor space. Higher density means more servers in a given area, usually accomplished by stacking equipment in racks.

drive logic — The circuitry included in the floppy or hard drive that interfaces with the disk controller.

EIA (Electronic Industries Alliance) unit (U) — A rack unit of measure equaling 1.75 vertical inches (4.45 cm).

generator kick (also **kick** or **kickstart**) — The time required for the generator to come online.

inrush power — Temporary surge of power to the server when it is turned on.

keyboard, video, mouse (KVM) — A console that enables you to control multiple servers from a single keyboard, video monitor, and mouse.

line conditioner — A device that filters out power inconsistencies, temporarily bridges power in the event of a brief brownout, suppresses high voltage spikes, and provides overall buffering between building power and the system.

load equipment — Anything connected to the UPS that draws power, usually servers and possibly other network equipment.

N+1 — A term that describes the expandability of a given server component or components, or space provided for expandable components. "N" is a variable that refers to the quantity of a given component installed in a system, and "+1" refers to a spare component.

power distribution unit (PDU) — A device similar in function to a household power strip that connects multiple devices to a power supply, but it is capable of much higher power capacity.

power factor — The difference between actual power and apparent power.

power supply unit (PSU) — The internal power supply powering a server or servers.

run time — The number of minutes that batteries can power the system.

standby power supply (SPS) — A device or technique that detects an interruption in line power and switches to a large transformer that stores a small amount of power required to bridge the time it takes to switch to battery power.

thermistor — A power supply thermostat that increases or decreases fan speed based on heat generated by the power supply.

Review Questions

1. Which of the following power supply wattages is likely to be found in a server?
   a. 145 W
   b. 200 W
   c. 250 W
   d. 375 W
2. How is the physical connection of a non-hot-swappable power supply different from a hot-swappable power supply? Choose two.
   a. You must remove individual power connectors.
   b. You must remove the server case cover.
   c. The PSU plugs directly into the chassis without individual cable connections.
   d. The wattage is much higher.
3. What is N+1 redundancy?
   a. It is a measure of the number of functioning hard disks plus one hot spare hard disk.
   b. "N" refers to the quantity of a given component and "+1" refers to a spare.
   c. It is a modem command that dials without waiting for a dial tone.
   d. It refers to the number of processors plus one redundant processor.
4. What might prevent you from turning on an entire rack of equipment at once?
   a. nothing, as long as a UPS is attached
   b. a UPS with not enough capacity
   c. inrush power
   d. a UPS with not enough run time

5. How many amps might a PCI slot use?
   a. 2
   b. 3
   c. 4
   d. 5
6. What is the load equipment?
   a. anything connected to the UPS that draws power
   b. relatively heavy rack equipment
   c. any device requiring more than 1 amp
   d. UPS batteries
7. What is the purpose of a UPS?
   a. to provide extended run time in the event of a power failure
   b. to allow the server to go into a low-power, energy-saving state
   c. to conserve electricity
   d. to provide enough power for a graceful shutdown
8. What is the advantage of the online/double conversion UPS?
   a. no switchover time
   b. reduced energy efficiency
   c. low cost
   d. only converts DC current to AC during a blackout or brownout
9. How does a UPS trigger UPS software to perform actions when power fails?
   a. using a USB or serial port to communicate with the software
   b. Nothing. The software detects dips in power or blackouts by itself.
   c. infrared signals
   d. a directed send through the network
10. What good is N+1 redundancy with a UPS?
    a. It provides additional run time in the event of a power failure.
    b. If a UPS fails, the other continues to provide power redundancy.
    c. It provides additional power capacity.
    d. N+1 is not a UPS function.
11. Where in the rack should you usually place UPS equipment?
    a. at the top
    b. in the middle
    c. anywhere
    d. at the bottom

12. Where are you likely to install additional fans in the rack?
    a. in one of the side panels
    b. in the bottom panel above the floor
    c. in the top panel
    d. in the back of the server
13. What is a primary advantage of a rack?
    a. increased density
    b. less heat buildup
    c. quieter operation because of the doors
    d. better organization of equipment
14. How big is 1U?
    a. depends on the rack manufacturer
    b. 1.75 inches (4.45 cm)
    c. 2.4 inches (6.07 cm)
    d. 1.57 inches (3.99 cm)
15. How much total approximate weight can a full height rack accommodate?
    a. 2000 lbs
    b. 1000 lbs
    c. limited only by floor support
    d. virtually unlimited
16. Where should you generally place the heaviest items in the rack?
    a. top
    b. middle
    c. bottom
    d. anywhere, as long as the overall balance is good
17. What helps to manage cables in the rack?
    a. cable management arms
    b. zip ties
    c. duct tape
    d. the PDU
18. What is the purpose of a KVM?
    a. to use a single console to control multiple computers
    b. to use multiple consoles to control a single computer
    c. to shorten the cable length of I/O devices
    d. to use the multi-monitor capabilities of Windows 2000

19. How does a PDU differ from a standard power strip?
    a. the number of outlets
    b. higher power capacity
    c. lower power capacity
    d. no difference
20. What is the purpose of a ballast?
    a. to add weight near the top of the rack
    b. to add weight to the side of a rack
    c. to fill empty spaces in the front of the rack
    d. to add weight near the bottom of the rack

HANDS-ON PROJECTS

Project 4-1

There are several things wrong with the rack depicted in Figure 4-26. Name them, and explain why they are wrong. On a separate piece of paper, diagram a more sensible configuration for this rack.

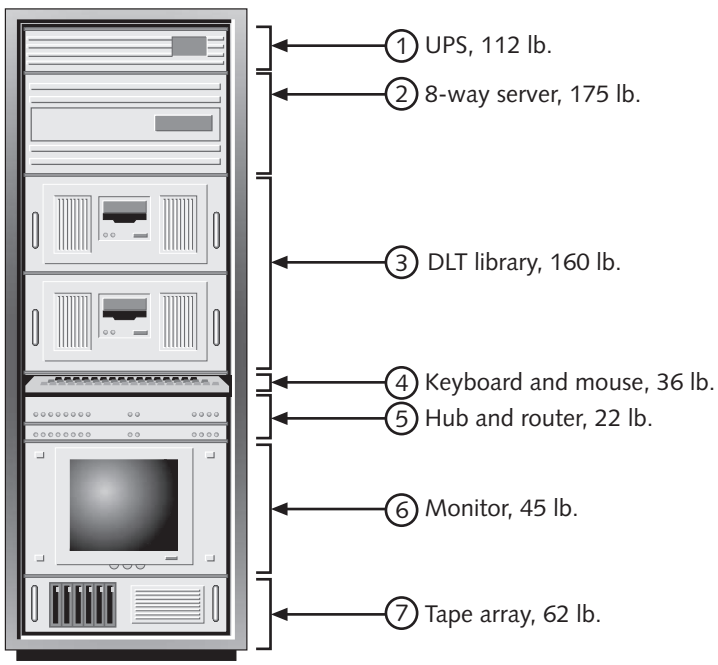

Figure 4-26 Server rack configuration for Project 4-1

Project 4-2

In this project, you will change a redundant power supply.

1. Ensure that server power is on with at least one redundant power supply.
2. Remove one of the power supplies. You might have to first remove the cover.
3. Notice that the power service to the server continues to function.
4. If the power redundancy is 2+1, and the remaining two power supplies provide load balancing (meaning that a single power supply can power the system), then remove another power supply. Power service to the server continues to function.
5. Leave the power supplies out for a few minutes. You might notice the fan on the last remaining power supply speeds up due to higher temperatures as it solely provides power.
6. Reinstall the power supplies you removed earlier, and put the cover back on if you removed it. Shut down the server for the next project.

Project 4-3

In this project, you will connect a UPS to the server.

1. Ensure that the server is powered off.
2. Plug the UPS power cable into utility power. Anything over 200 V might be a locking plug, which prevents accidental removal.
3. Connect the server's power cord to the UPS.
4. Connect the UPS serial or USB cable to the server. Note that the serial cable is usually proprietary to the UPS — you cannot use a standard serial cable.
5. Turn on the server.
6. Verify that the battery is installed (or connected, if external), and then turn on the UPS. Depending on the brand, this might involve simply pushing the power button or pushing and holding it for a few seconds. The battery is probably not fully charged (it usually takes several hours to fully charge a battery).
7. Observe the indicators on the UPS. You can probably see indicators that show the charge percentage, load level, and so forth.
8. When the NOS has loaded, log on and install the UPS management software (such as APC PowerChute).
9. Using the software, access the interface that allows you to configure email or pager alerts when a UPS event (such as power failure) happens. Enter your email address and/or pager number. If connected to the network or phone line, you should receive notification with each configured UPS event.
10. Configure the software to automatically save open documents.
11. Configure the software to automatically shut down after two minutes (you might want more time in the real world).

12. Open a document in a word processor, type in some words, and save the document. Then, type in some more words without saving the document.
13. If the battery has at least a 50% charge, unplug utility power.
14. Watch the system shut down the server.
15. Plug utility power back in, and turn on the server.
16. When the NOS loads and you log on, open the document you created in Step 12. Notice the changes were saved for you automatically.
17. Later, check your email and/or pager for a notification of the failed utility power.
18. Optionally, restart the server but do not let the OS load. You can do this by inserting a blank floppy in the floppy drive assuming the BIOS is set to boot from the floppy. Note the battery level, and then unplug utility power and see how many minutes and seconds the battery is able to power the system.

Project 4-4

In this project, you will choose a UPS system using the APC web site.

1. Using your web browser, go to *www.apc.com*.
2. Navigate to the Selectors section, and choose **Go to the UPS Selector**.
3. Notice that you can select several types of electronic equipment for a UPS. Click on **Server**.
4. You are presented with a list of various servers. Choose a high-end PC server. Suggestions: Compaq 8-way, Dell PowerEdge 8450 rack mount, HP Netserver LXr 8500dc.
5. Enter information about the server such as the chassis type (select **Rack Mount**), monitor, number of processors, and so forth. Enter whatever you like, but be somewhat realistic; otherwise, there won't be a UPS powerful enough. Because the voltages are likely to be high, select a NEMA L6-30P plug type (a locking, 30 amp plug).
6. You are returned to the initial screen to choose more equipment if necessary. You can add more if you like, and when finished, continue on to the Preferences section.
7. You might have to back up and make a change here or there, but eventually you will find one or more solutions from APC.
8. On a separate sheet of paper, answer the following questions:
   a. How long will the server run with this UPS?
   b. What percent of the maximum capacity is used by this UPS?
   c. How many watts?
   d. What type of connections (plugs and receptacles) does it offer?
9. Close the web browser.

Project 4-5

In this project, you will configure a rack using HP Netserver Rack Assistant.

1. Navigate to the Hewlett-Packard web site at *www.hp.com*.
2. Access the **Products and Services** section of the site.
3. Under Server, select **PC Server**.
4. Under Accessories, select **Rack Solutions**.
5. Navigate to the link for the rack assistant. (Try a search if you cannot navigate your way to the link.)
6. Download the software and install it on a Windows computer.
7. Start the software, and open a new rack when prompted.
8. Choose the rack attributes. For the rack voltage, select 200/208. Leave the rules selected. The software offers you suggestions about your configuration and stack equipment according to predefined rules. Use a stand-alone rack.
9. Add items to the rack as you like by double-clicking them. Add at least two servers, two mass storage units, a keyboard/monitor, and at least one UPS. You will notice that as you add items, statistics about the weight, voltage, BTU/hour (heat measurement), and so forth appear on the right.
10. When you print or save the document, the Rack Assistant offers you suggestions that you can accept or deny.
11. Close the Rack Assistant, and then close the web browser.

Project 4-6

In this project, you will install rack equipment in the rack.

1. Access a rack at least 24U in height. Have various pieces of equipment that you can add to the server, such as at least one server, KVM, monitor, PDU, and UPS. If available, also try to have cable management equipment.
2. Arrange the equipment from heaviest to lightest. This is not necessarily the order in which you will install the equipment, but it is a guideline.
3. If the rack has stabilizing feet (highly recommended), extend or attach them.
4. Install the UPS in the bottom of the server. Remember that to install the UPS and all other heavy items, you should get help from another person.
5. Install one or more PDUs. Plug the PDU into the UPS. Do not plug in the UPS. This project is for installation purposes only; you will not be powering up equipment.
6. Install either the server or the mass storage device above the UPS. This can go either way depending upon the weight of the mass storage device. Lighter storage devices can be installed higher in the rack or where it is easier to access the equipment to change tapes.
7. Place the KVM at a comfortable height and the monitor at a viewable height.

8. Route the cables along cable arms and brackets if available.
9. If you have two or more servers to install, use the KVM to make appropriate connections between the two. Remember that for each server, you need one set of extension cords for the keyboard and mouse (both round PS/2 connectors) and a video cable. If you have two servers, then you need two sets of extension cords.

Case Projects

1. Your multimedia organization has recently absorbed a small corporation in Hutchinson, KS. Your supervisor wants you to travel to Hutchinson to ensure that the single server in that office is as close as possible to the same hardware reliability as servers in the main office. You arrive at the new office and find that the "server" is really just a high-end PC workstation with Windows NT 4.0 installed on it. In addition to its role as a server, the graphic artists also access the computer for high-end graphics work. You are concerned about all the modifications to the system. On a 250 W power supply, the server has the following installed hardware running 24/7:

Component	Watts
Motherboard	30 W
Pentium III 750 MHz	25 W
4X AGP	30 W
256 MB RAM	20 W
Four occupied PCI slots	20 W
3.5 inch floppy drive	5 W
3 IDE 7200 internal hard disks	45 W
2 CD-ROMs	50 W
Total wattage	**225 W**

   What problem(s) do you see with this server and what solution(s) would you recommend?

2. Your large enterprise has grown very quickly over the past three years. You have experienced difficulty keeping up with rapid expansion, but you are traveling to several offices over the next few months to try to ensure that each site has the necessary power redundancy available. You visit the Livermore, CA site and immediately hear concern from the local administrator that the last time a brownout occurred, the UPS system in one of the racks provided power for only five minutes before losing power. You analyze the rack in which the UPS is installed, and it seems that the UPS carries sufficient VA to power the rack for at least 15 minutes. What could be causing the short power duration?